MATHia X: The Next Generation Cognitive Tutor

Steven Ritter
Stephen E. Fancsali
Carnegie Learning, Inc.
437 Grant Street, 20th Floor
Pittsburgh, PA 15219, USA
1.888.851.7094 {x122, x219}
{sritter, sfancsali}
@carnegielearning.com

ABSTRACT

MATHia X is the next generation implementation of Carnegie Learning's Cognitive Tutor (CT), a widely deployed, research-based mathematics curriculum that has provided data for many educational data mining studies. While many researchers are familiar with the basic operation of the system, there are several features that may affect analysis and interpretation of data that are less well known. We describe features of MATHia X and CT, as well as aspects of its practical implementation in real-world classrooms, that may be important for researchers using MATHia X and CT datasets.

Keywords

MATHia X, Cognitive Tutor, intelligent tutoring systems, real-world implementation, mastery learning, wheel-spinning

1. MATHIA X & COGNITIVE TUTOR

MATHia X is the next generation platform for Carnegie Learning's Cognitive Tutor (CT) [5], an intelligent tutoring system (ITS) for mathematics used by hundreds of thousands of learners in middle schools, high schools, and universities across the US (and to a lesser extent internationally, e.g., [4]).

MATHia X provides an HTML5/JavaScript, web-based implementation of the Cognitive Tutor technology and mathematics curricula; for our mid-2016 release we will have content for middle school grades 6-8 and Algebra I, with subsequent content covering Algebra II and Geometry. While MATHia X provides a technology and user interface refresh (including a space-themed interface "skin" in the initial release), fundamentally, most user interface and ITS affordances (including fine-grained data collected about learner interactions in the ITS) are essentially the same as they were in the Java-based Cognitive Tutor and MATHia products that have been in use for well over a decade. As such, we expect to continue in our long-standing tradition of partnering with education, educational data mining, and cognitive science researchers on basic and applied research about how students think and learn, as well as to continue providing data to these communities. The present demo explains a number of features common to both our legacy CT product as well as our next generation MATHia X product, many of which are important to data analyses carried by educational data mining researchers.

Datasets from CT are widely used in a variety of educational data mining (EDM) and education research projects, including in a substantial number of papers in the proceedings of the present conference. Many experimental and observational datasets (comprising hundreds of millions of learner actions in CT) have also been made available via the Pittsburgh Science of Learning Center's DataShop repository [3]. While many aspects of MATHia X and CT, such as their use of mastery learning and Bayesian Knowledge Tracing (BKT) are well known, there are many features and details of implementation and context of use that are less well known but important for appropriate analysis of CT (and eventually MATHia X) data. We describe a number of these characteristics here, in the hope that this information can inform EDM researchers' understanding of CT and MATHia X and contribute to future research that uses such data.

2. FEATURES & IMPLEMENTATION

2.1 Basal and Supplementary Use

Carnegie Learning produces text materials in addition to software, and the "blended" product (text and software) is often used as a "basal" curriculum, meaning that it is the primary source of instructional materials for a class. Our recommendation for blended implementations is that the software be used approximately 40% of the time (two class periods/week), with the text materials used for 60% of classroom time. Depending on school schedules, computer availability, and other factors, the amount of software usage varies considerably between schools.

In addition to "basal" usage, some schools use CT as a supplement to other educational materials. Such usage may follow the 60%-40% model, using a different textbook, but most supplemental usage is irregular. One consequence of such usage is that estimates of student knowledge can be highly inaccurate, since students may learn (or forget) substantial amounts in the long gaps between use of the tutor. Some supplementary use is for a specific purpose (e.g., summer school). In both types of implementations, schools may use the software for all students or for only a subpopulation thereof (e.g., those below grade level).

2.2 (Custom) Curricular Structure

Within K-12, there are a variety of main Carnegie Learning curricula: Algebra 1, Geometry and Algebra 2 (the high school sequence) are provided by our legacy CT product; a three-year middle school sequence and Algebra I are provided by the new MATHia X product in its initial release; and Bridge to Algebra, a one-year review of the middle school sequence is also provided on our legacy platform. Soon all of our curricula will be provided on the web-based technology that drives MATHia X. Overall, these curricula correspond to typical US courses. However, depending on state standards and other needs, schools may construct "custom" curricula that incorporate topics from one or more of these prototypical curricula. Custom curricula are popular, and the majority of CT data is now collected within such custom sequences. CT validates custom sequences for redundancies and violations of prerequisites; schools can ignore warnings about violations, but this is rare.

A curriculum consists of a set of modules, which represents a major topic in the curriculum. A full course may contain 6-8 modules. Modules consist of units, which consist of sections. Each section contains a large set of problems. Mastery learning operates at the section level; students work within a section until they have mastered all associated knowledge components (KCs) (i.e., skills). The next section (or unit, if the section mastered is the final one in the unit) is automatically presented to the student. The module level is different. Although students will automatically progress to the next module when they complete the final section in the prior module, teachers can also "unlock" modules, allowing students to work on any open modules. Thus, at any given time, a student has a single position within a module (representing the current section) but may have positions within multiple modules. This feature is intended to allow movement among topics that do not have a prerequisite relationship.

2.3 Violations of Mastery Learning

Although we say CT and MATHia X implement mastery learning, in practice, there are several cases where students are not asked to work until they complete with mastery. Within each section of a curriculum, we specify a maximum number of problems that will be presented to students (often 25, but this varies, depending on the complexity of problems; for technical reasons, there are also cases where students might be promoted before reaching this maximum). If students complete this maximum without mastering their skills, they will advance to the next section of the curriculum. We call these advances "promotion," and these are flagged and communicated to teachers in our reporting system. The underlying idea is similar to the concept of "wheel-spinning" [1]. If students are not able to master the material in the tutor in a reasonable period of time, then it is likely that, for whatever reason, the tutor's mode of instruction for this topic is not resonating with the student, and so an alternate instructional approach is preferable. The teacher is responsible for presenting the alternative approach. Promotion is not rare; students are promoted from about 12% of sections. Promotions vary quite a bit by section and by student. Teachers also have the ability to manually move a student to a different position in the curriculum. Such placement changes also violate the mastery assumption. They happen for various reasons, most commonly because the teacher wants the student to "catch up" to the placement of the rest of the class. Such mastery learning violations due to placement changes are associated with greater error rates (and greater variability in error rates) over time than those experienced by students in classes that do not violate mastery learning [6].

2.4 Instructional Resources

Many analyses of CT data have looked at help seeking (e.g., [7]). Such work typically considers student use of problem-specific help, which is the only resource that affects CT's assessment of student knowledge. However, there are other sources of assistance available. Each unit has "lesson" content, which provides declarative instruction, worked examples, manipulatives, and topic-related video. A glossary is always available to students, and references to math terms within lesson text or hints are linked to it. Students also often use calculators and communicate with teachers and other students as they use the software.

Step-by-step examples provide another form of assistance. At least one example problem in each unit illustrates the basic problem-solving approach [2]. Unlike "regular" problems, step-by-step examples expose only one possible path through the problem, and text that would be used as a hint in problem solving is automatically presented to students as they go through the step-by-step example. This experience is intermediate between looking at a worked example and problem solving. Students can refer back to the step-by-step example as they work, and work in the step-by-step example is not used to assess student knowledge.

2.5 Non-persistent Student Model

Math knowledge is cumulative, so one expects that new topics incorporate many KCs mastered in earlier topics. Each section in CT and MATHia X monitors a small set of KCs, among the large set that is actually needed to solve problems in the section. While each section does introduce new knowledge, for various reasons, some sections list KCs that have been addressed in previous sections. These KCs take their preset values, not values based on students' prior work. In other words, CT and MATHia X do not assume that such KCs have been mastered. There is little practical consequence to listing such KCs; if the student learned them, CT will quickly recognize that fact, but researchers should be aware that the CT's assessment of skills is always within a section. Since skill values (i.e., estimates of student knowledge of a skill) do not carry over from section to section, researchers should not automatically assume that KCs with identical names in different sections are, in fact, identical KCs for purposes of data analysis.

3. DEMO + THE FUTURE

In this demo, we will exhibit basic problem solving in MATHia X, introducing the Cognitive Tutor technology to those unfamiliar with it and showing the refreshed technology to those already familiar with our products. Carnegie Learning looks forward to broad adoption of the next generation MATHia X software as a part of its blended mathematics curricula. Combining observational data sets from such adoptions with experimental data sets that will be collected by investigators using MATHia X as a platform for research will provide rich data to be mining and analyzed for many years to come in the educational data mining, learning analytics, cognitive science, and other research communities.

4. REFERENCES

[1] Beck, J.E., Gong, Y. 2013. Wheel-spinning: Students who fail to master a skill. In *Proceedings of AIED 2013* (Memphis, TN, Jul. 2013), 431-440.

[2] Hausmann, R.G.M., Ritter, S., Towle, B., Murray, R.C., Connelly, J. 2010. Incorporating interactive examples into the Cognitive Tutor. In *Proceedings of ITS 2010* (Pittsburgh, PA, Jun. 2010), 446.

[3] Koedinger, K.R., Baker, R.S.J.d., Cunningham, K., Skogsholm, A., Leber, B., Stamper, J. 2011. A data repository for the EDM community: the PSLC DataShop. In *Handbook of Educational Data Mining*, C. Romero, S. Ventura, M. Pechenizkiy, & R.S.J.d. Baker, Eds. CRC, Boca Raton, FL.

[4] Ogan, A., Walker, E., Baker, R., Rebolledo, G., Jimenez-Castro, M. 2012. Collaboration in Cognitive Tutor use in Latin America: Field study and design recommendations. In *Proceedings of CHI 2012* (Austin, TX, May 2012), 39-48.

[5] Ritter, S., Anderson, J.R., Koedinger, K.R., Corbett, A.T. 2007. Cognitive Tutor: applied research in mathematics education. *Psychon. B. Rev.* 14 (2007), 249-255.

[6] Ritter, S., Yudelson, M.V., Fancsali, S.E., Berman, S.R. 2016. How Mastery Learning Works at Scale. In *Proceedings of the 3rd Annual ACM Conference on Learning at Scale* (Edinburgh, Scotland).

[7] Roll, I., Aleven, V., McLaren, B.M., Koedinger, K.R. 2011. Improving students' help-seeking skills using metacognitive feedback in an intelligent tutoring system. *Learn. Instr.* 21 (Apr. 2011), 267-280.